

The History Heuristic and Alpha-Beta Search Enhancements in Practice

by Jonathan Schaeffer

Tsan-sheng Hsu

徐讚昇

tshsu@iis.sinica.edu.tw

<http://www.iis.sinica.edu.tw/~tshsu>

Abstract

- Introduce a new move ordering technique called **history heuristic** for better move ordering.
- Study the effect of combining multiple heuristics.
 - Each enhancement should not be taken in isolation.
 - Try to find the combination that provides the greatest reduction in tree size.
- Be careful on artificial game trees.
- Be careful on the type of game trees that you do experiments on.
 - Depth, width and leaf-node evaluation time.
 - A heuristic that is good on the current experiment setup may not be good some years in the future because of the the game tree can be evaluated much deeper in the the same time using faster CPU's.

Commonly used heuristics

- Iterative deepening
- Transposition table
- Minimal window search: Scout or NegaScout
- Aspiration search
- Killer heuristic
- Knowledge heuristic
- New heuristic: the history heuristic

Intuition

- The size of the search tree built by a depth-first alpha-beta search largely depends on the order in which branches are considered at interior nodes.
- It looks good if one can search the best possible subtree first in each interior node.
- A better move ordering normally means a better way to prune a tree by the alpha-beta search.
- Enhancements to the alpha-beta search have been proposed based on one or more of the following principles:
 - move ordering;
 - window size;
 - re-using information.

Transposition tables

- We are searching a game graph, not a game tree.
 - Interior nodes of game trees are not necessarily distinct.
 - It may be possible to reach the same position by more than one path.
- What's in an entry of a **transposition table**?
 - The position p .
 - Searching depth d .
 - Best value in this subtree.
 - Best move for this position.
- How to use information in the transposition table?
 - Suppose p is searched again with the depth limit d' .
 - If $d \geq d'$, then no need to search anymore.
 - ▷ *Just retrieve the result from the table.*
 - If $d < d'$, then use the best move stored as the starting point for searching.
- Need to be able to find p in a large table efficiently.

Zobrist's hash function

- Find a hash function $hash(p)$ so that with a very high probability that two positions will be mapped into distinct locations in the table.
- Using XOR to achieve fast computation:
 - associativity: $x \text{ XOR } (y \text{ XOR } z) = (x \text{ XOR } y) \text{ XOR } z$
 - commutativity: $x \text{ XOR } y = y \text{ XOR } x$
 - $x \text{ XOR } x = 0$
 - ▷ $x \text{ XOR } 0 = x$
 - ▷ $(x \text{ XOR } y) \text{ XOR } y = x \text{ XOR } (y \text{ XOR } y) = x \text{ XOR } 0 = x$
 - $x \text{ XOR } y$ is random if x and y are also random

Hash function

- Assume there are k different pieces and each piece can be placed into r different locations.
 - Obtain $k \cdot r$ random numbers in the form of $s[piece][location]$
 - $hash(p) = s[p_1][l_1] \text{ XOR } \dots \text{ XOR } s[p_x][l_x]$ where p_i is the i th piece and l_i is the location of p_i .
- This value can be computed incrementally.
 - Assume the original hash value is h .
 - A piece p_{x+1} is placed at location l_{x+1} , then
 - ▷ *new hash value = $h \text{ XOR } s[p_{x+1}][l_{x+1}]$.*
 - A piece p_y is removed from location l_y , then
 - ▷ *new hash value = $h \text{ XOR } s[p_y][l_y]$.*
 - A piece p_y is moved from location l_y to location l'_y then
 - ▷ *new hash value = $h \text{ XOR } s[p_y][l_y] \text{ XOR } s[p_y][l'_y]$.*

Clustering of errors

- **Though the hash codes are uniformly distributed, the idiosyncrasies of a particular problem may produce an unusual number of clashes.**
 - **if $hash(p^*) = hash(p^+)$, then**
 - ▷ *adding the same pieces at the same locations to positions p^* and p^+ produce the same clashes;*
 - ▷ *removing the same pieces at the same locations from positions p^* and p^+ produce the same clashes.*

Practical issues

- Normally, design a hash table of 2^n entries, but with key length $n + m$ bits.
 - That is, each $s[*piece*][*location*]$ is a hash value of $n + m$ bits.
 - Hash index = $hash[p] \bmod 2^n$.
 - Store the hash key to compare when there is a hash hit.
- How to store a hash entry:
 - Store it when the entry is empty.
 - Replace the old entry if the current result comes from a deeper subtree.
- Errors:
 - Assume this hash function is uniformly distributed.
 - The chance of error for hash clash is $\frac{1}{2^{n+m}}$.
 - Assume during searching, 2^w nodes are visited.
 - The chance of no clash in these 2^w visits is $P = (1 - \frac{1}{2^{n+m}})^{2^w} \simeq (\frac{1}{e})^{2^{n+m-w}}$.
 - When $n + m - w$ is 10, $P = 0.99901$.
 - Currently:
 - ▷ $n + m = 64$
 - ▷ $w \leq 32$.

Refutation tables

- For each iteration, the search yields a path for each move from the root to a leaf node that results in either the correct minimax value or an upper bound on its value.
 - This path is often called **principle variation** or **principle continuation**.
 - Store the current best principle variation at P_i for each depth i .
- This path from the $d - 1$ ply search can be used as the basis for the search to d ply.
 - Assume using iterative deepening.
- Searching the previous iteration's path or **refutation** for a move as the initial path examined for the current iteration will prove sufficient to refute the move one ply deeper.
 - When searching a new node at depth i ,
 - ▷ *try the moves made by this player at P_i first starting from beginning to the end;*
 - ▷ *then try moves made by this player at P_{i+1} starting from beginning to the end;*
 - ▷ ...

Aspiration search

- The normal alpha-beta search usually starts with a $(-\infty, \infty)$ search window.
- If some idea of the range of the search will fall is available, then tighter bounds can be placed on the initial window.
 - The tighter the bound, the faster the search.
 - Some possible guesses:
 - ▷ *During iterative deepening, assume the previous best value is x , then use $(x - \text{threshold}, x + \text{threshold})$ as the initial window size where threshold is a small value.*
- If the value falls within the window then the original window is adequate.
- Otherwise, one must re-search with a wider window depending on whether it fails high or fails low.
- Reported to be at least 15% faster than the original alpha-beta search.

Killer heuristic

- A compact refutation table.
- Storing at each depth of search the moves which seem to be causing the most cutoffs, i.e., so called **killers**.
 - Currently, store two most recent cutoffs at this depth.
- The next time the same depth in the tree is reached, the killer move is retrieved and used, if valid in the current position.

History heuristic

■ Intuition:

- A move M may be shown to be best in one position.
- Later on in the search tree a **similar** position may occur, perhaps only differing in the location of one piece.
 - ▷ *A position p and a position p' obtained from p by making one or two moves are likely to share important features.*
- Minor difference between p and p' may not change the position enough to alter move M from still being best.

■ In alpha-beta search, a **sufficient**, or good, move at an interior node is defined as one that

- causes a cutoff, or
- if no cutoff occurs, the one yielding the best minimax score.

Implementation I

- Keep track of the history on which move is a good move.
 - Assume the board has q different locations.
 - Assume each time only a piece can be moved.
 - There are only q^2 possible moves.
 - Including more context information, e.g., the piece moving, did not significantly increase performance.
 - ▷ *If you carry the idea of including context to the extreme, the result is a transposition table.*

Implementation II

- Each time when a move is a good move, increases its counter by a certain **weight**.
 - During move generation, pick one with the largest counter value.
 - ▷ *Need to access the history table and then sort the weights in the move queue.*
 - The deeper the subtree searched, the more reliable the minimax value except in pathological trees, rarely seen in practice.
 - The deeper the search tree, and hence larger, the greater the differences between two arbitrary positions in the tree and less they may have in common.
 - By experiment: let $\text{weight} = 2^{\text{depth}}$, where *depth* is the depth of the subtree searched.
 - ▷ *Several other weights, such as 1 and depth, were tried and found to be experimentally inferior to 2^{depth} .*
- Killer heuristic is a special case of the history heuristic.
 - Killer heuristic only keeps track of one or two successful moves per depth of search.
 - History heuristic maintains good moves for all depths.
- History heuristic is very **dynamic**.

Experiments: Setup

- **Try out all possible combinations of heuristics.**
 - 6 parameters with 64 different combinations.
- **Searching depth from 2 to 5 for all combinations.**
 - Applying searching upto the depth of 6 to 8 when a combination showed significant reductions in search depth of 5.
- **A total of 2000 VAX11/780 equivalent hours are spent to perform the experiments.**

Experiments: Results

■ Results:

● Using a single parameter:

- ▷ *History heuristic performs well, but its efficiency appears to drop after depth 7.*
- ▷ *Knowledge heuristic adds an additional 5% time, but performs about the same with the history heuristic.*
- ▷ *The effectiveness of transposition tables increases with search depth.*
- ▷ *Refutation tables provide constant performance, regardless of depth, and appear to be worse than transposition tables.*
- ▷ *Aspiration and minimal window search provide small benefits.*

● Using two parameters

- ▷ *Transposition tables plus history heuristic provide the best combination.*

● Combining three or more heuristics do not provide extra benefits.

Comments

- **Combining two best heuristics may not give you the best.**
- **Need to weight the amount of time spent in realizing a heuristic and the benefits it can bring.**
- **Need to be very careful in setting up the experiments.**

References and further readings

- * J. Schaeffer. The history heuristic and alpha-beta search enhancements in practice. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 11(11):1203–1212, 1989.
- * A. L. Zobrist. A new hashing method with applications for game playing. Technical Report 88, Department of Computer Science, University of Wisconsin, Madison, USA, 1970. Also in *ICCA journal*, vol. 13, No. 2, pp. 69–73, 1990.
- * Selim G. Akl and Monroe M. Newborn. The principal continuation and the killer heuristic. In *ACM '77: Proceedings of the 1977 annual conference*, pages 466–473, New York, NY, USA, 1977. ACM Press.
- S.C. Hsu. Searching Techniques of Computer Game Playing. *Bulletin of the College of Engineering, National Taiwan University*, 51:17–31, 1991.